STANFORD UNIVERSITY MEDICAL EXPERIMENTAL COMPUTER RESOURCE

S U M E X

Proposal for a research resource in extension of the ACME project.

A. INTRODUCTION

1. Objectives and Long Term Goals

General purpose computer support of research at Stanford University Medical School has reached substantial maturity under the impetus of the ACME project (Advanced Computer for Medical Research) funded since 1966 by NIH. We have understood that our technical success would be coupled with the gradual withdrawal of centralized agency support for a proven utility. Hence, June 1973 represents the termination of the longstanding NIH subsidy for ACME, which will thenceforth be operated as a fee-based service.

The present application seeks to establish a resource for a new set of technical horizons, in keeping with the expending capability and applications of computers in biomedical research. SUMEX would be a resource (1) to support a set of ongoing biomedical research programs that exploit state-of-the-art computer techniques, and help to shape further advances, and (2) the computer-science research that is essential to expedite the creative uses of computers in the laboratory, in the clinical research wards, and eventually in patient care at every level.

The unifying theme of the SUMEX resource is the management of a set of peripheral minicomputers by a powerful central facility. The minicomputers in question are immediately situated in laboratory and clinical-research environments. They can perform some tasks free-standing, but the selected projects require further backup to sustain high-data-rate and closed-loop operations. These machines, together with others on less demanding projects, can be enhanced by sharing peripherals, mutual backup, and higher level language programming and debugging in the central processor. More far-reaching in concept, and central to this proposal, is central and "intelligent" management of the data-gathering process to meet problem-oriented needs for information. (This is no more than a feeble emulation of the processes that higher organisms must have evolved to modulate the flow of sensory data into the perceptual mechanism.)

In the ideal situation this might entail a realtime closed-loop control of a laboratory instrument or a patient-monitoring device. Prior and currently updated information, related to partial solutions of a problem, would then selectively orient further data-taking so as to expedite a complete solution. A related example would be a kind of triage--allocating the time-shared partitions of a large computer resource to concentrate on the patient with the most problematical symptoms. Even where realtime processing is unrealistic, in the present state of the art, as in motion-picture processing, the magnitude of computing requirements could be vastly reduced by analyzing each frame to pose specific questions of the next one, rather than prepare a digital core image of the entire sequence. Similar problems arise in every branch of spectrometry, including mass-spectrometry, where costly instruments and samples may be needlessly expended in conventionally serial acquisition of the whole spectrum, followed by its analysis (which usually relies upon a small part of the entire data set.)

The initial list of collaborating investigators is presented in part C-1. Others who are not yet prepared to commit themselves to this enterprise will continue to be recruited as discussed in part D. The hardware requirements are detailed in part E and further details of the operation of the SUMEX resource, and its relationship to service-computing at SUMC* are detailed in subsequent sections.

*(Stanford University Medical Center)

2. Background: ACME Computer Facility Experience

On August 1, 1966, the Biotechnology Research Resources Branch of NIH (then known as the Research Resources Branch) awarded a grant to Stanford University Medical School to support the establishment of ACME (Advanced Computer for MEdical Research) facility. The initial proposal included the following paragraph concerning hardware selection and resource allocation:

> "The IBM/360-50 has been selected for the initial realization of ACME (1) as a machine technically appropriate to the immediate tasks in mind and (2) for its system compatibility with the 360-67 already selected for the eventual replacement of the 7090 by the Stanford Computation Center. The 360-50 will be installed in ACME May 1966 and will run on three shifts under Operating System/360, subject to review by the policy committee. These will be dedicated respectively:
>
> (A) A prompt access time-sharing mode - perhaps over most of the working day.
>
> (B) A scheduled, full-use, on line mode - to service development work on high data rate and on line control applications, and for similar systems development.
>
> (C) Job-shop, especially longer runs for which overnight turnaround is acceptable, and which cannot be serviced with comparable effectiveness by SCC."

The following aims were added to the ACME charter at the time of the Renewal Proposal in the Spring of 1969:

(a) To improve hardware and software reliability for the benefit of the medical users.

(b) To provide small machine assemblers in PL/ACME so that code for small machines can be written from an ACME terminal.

(c) To achieve over time a state where income from user charges will match operational costs for the ACME system. The target date for this has not yet been fixed by Stanford and NIH.

All of the original objectives have been achieved to varying degrees of satisfaction. Of special note is the development of PL/ACME as an interactive timesharing system which can be easily learned and used by medical staff. On the other hand, the realtime support offered is inadequate due to system instabilities and data rate limitations. Access to Campus facility is inconvenient for ACME users.

In terms of the items added at renewal time, hardware and software reliability have been markedly improved. Small machine assemblers have been added, but the user must write code in the assembly code for whatever satellite he intends to run. At present, assemblers of this type exist for PDP-11, PDP-8, and 1800. The income of the facility has been rising steadily. Economic overlaps with NIH direct support for ACME have blurred the transition to totally non-subsidized use. A major rate increase was initiated in April, 1972. With this change, income in the near term is expected to reach roughly 60% of direct operating costs (exclusive of development efforts). From the vantage point of hindsight one could well ask whether the selection of the 360/50 hardware and the decision to promote a large central time-sharing and data collection resource were appropriate. Given the availability of new third generation hardware and the promises of IBM or expectations of its customers in 1966, the 360/50 hardware selection is defensible. However, the development of low cost, fast, well-supported minicomputers was not anticipated to proceed at the phenomenal pace that it has. This major technological shift has strongly influenced our present thinking for the future of computing in medicine and related research. The role of a large shared resource has by no means been obviated by the minicomputer revolution. We will continue to need powerful facilities beyond the scope of current mini architecture.

The advantages of dedicated satellite processors make them mandatory for many applications which require high reliability and availability. A marriage of of the two architectures is proposed. The resultant synergism is designed to solve identified problems in our research environment.

B. SUMMARY

The resource for which we are applying consists essentially of a Digital Equipment Corporation PDP-10 system, to be procured on a lease-purchase plan over the 5-year term of the grant. The configuration has been designed for flexibility in interfacing with numerous small machines)provided by the collaborating investigators). The rationale for this system, the associated peripherals, and for the technical support staff, are detailed herein below. The SUMEX machine will be scheduled for the sole use of the SUMEX collaborative group, under the leadership of the principal investigator. Other investigators will, however, be recruited into the group if they qualify by virtue of their interest in and competence for computer research relevant to the main themes outlined here, and insofar as their theoretical and technical contributions will enlarge our understanding of the application of computers for the management of high-data-rate biomedical studies. SUMEX will not be available for routine computing that can be effectively purchased from existing utilities. The initial group of coinvestigators includes all respondents at SUMC who qualified by the stated criteria; however, several others are expected to advance the sophistication of their computer applications so as to qualify within the term of the grant.

SUMEX is expected to develop a number of applications that may become routine once perfected. These results would be transferred as appropriate to a service utility which continues to meet the demand for conversational time-shared computing as a legacy of the ACME project. For example, a number of workers at SUMC have utility-level requirements for support of their minicomputers. We believe this cognate requirement can be most efficaciously met if SUMC establishes a twin <u>service</u> facility based on a PDP-10, perhaps coupled with a small IBM machine for fiscal data-processing with company-furnished software. (These decisions are outside the policy cognizance of the present applicants, and outside the funding hereby requested. However, we are in good communication with the SUMC computer service committee that is establishing those policies.)

These decisions need not be fixed at an early date; for example, the existing IBM/360-50 ACME system might be retained for some time to provide feepaid timesharing service, during a transition period at the establishment of SUMEX as a research resource. In any case, the existing ACME project, for the remainder of its term (Expires July 31, 1973) and to some extent the new SUMEX program will have the responsibility of easing that transition for the community of users who have made large commitments to the ACME service. Except during an interim transition period and later on an emergency basis, SUMEX obligations to the ACME community will be confined to providing technical advice for conversion, and developing software that can be used interchangeably on the SUMEX and on the SUMC service machines. However, SUMEX developmental efforts will be strongly biased by that requirement for compatibility, and indeed for ready exportability to other biomedical computer groups.

The application of SUMEX could be summarized in terms of the scientific objectives of specific projects -- which are, however, detailed in section G. These lend substance to the technical research on computers themselves which is the unifying theme of our proposal. The principal investigator's interest in the DENDRAL project was, from the outset, motivated by the aim of broadening the application of machine intelligence to science generally (He is, after all,a geneticist first, and entered into mass spectrometry only because the latter was more amenable at the present level of the art.). While few other fields of biomedical research are, at present, ready for the full-blown application of the techniques developed for DENDRAL as a problem in artifical intelligence, there is a broader base of common concern for related problems in data management. Briefly, in as many ways as possible, we will be developing the means to support small machines by a large central facility acting as an executive manager for the minicomputers. In addition, we will develop the technology of programming the mini's in higher level languages compiled on the PDP-10; will simulate minicomputer configurations as a way of designing new installations; will provide buffering and communications among small machines, and between them and various peripherals, including secondary storage, displays, and (if the opportunity materializes) access to other nodes on national computer networks. Where control loops are to be closed back on the mini's, a great deal of processing of the experimental data will presumably be done in the PDP-10. We will also investigate the utility of small machines as auxiliary subroutine-processors to increase the efficiency of a time-shared central device in some long computations.

C. JUSTIFICATION

1. Demonstration of Need

The recruitment of an initial group of collaborators has made clear the need for:

- dedicated large computing resources
- high data rate acquisition and control capabilities
- development of software and hardware techniques
- integration of host and satellite computing systems

The SUMEX Resource responds to needs identified by a number of research projects within the Medical Center. Individually, their projects are unable to avail themselves of resources as large as SUMEX. Collectively, their research objectives demand the capabilities designed into SUMEX in this proposal. The requested facility calls for a high degree of cooperation among a small number of collaborators for their mutual benefit.

Five collaborative projects are described as part of this proposal. They are:

a. Predictive modelling of cardiovascular function utilizing X-ray and ultrasonic imaging techniques, Dr. Donald C. Harrison, Cardiology.

b. Development of computer based characterizations of radiographs of ureter, Drs. Thomas Stamey and Chris Constantinou, Urology.

c. DENDRAL -- Computer automation of the interpretation of mass spectrum, Dr. Joshua Lederberg, Genetics; Dr. Carl Djerassi, Chemistry; and Dr. Edward Feigenbaum, Computer Science.

d. Cell separator automation, Drs. L. Herzenberg and E. Levinthal, Genetics.

e. Electroencephalogram Driven Stimulus/Response Studies of Drug Effects, Drs. B. Kopell, T. Roth, and Pfefferbaum, Psychiatry.

2. SUMEX Relationship to Institutional Plans.

Stanford University has relied extensively on IBM equipment; the decision to procure a PDP-10 for SUMEX inevitably hinders compatibility and economy in system effort. Unfortunately, the IBM lines offer no cost-effective equivalent to the PDP-10 for small machine interfacing. An oversized and costly IBM machine would offer few advantages of portability to other research programs interested in similar objectives. Furthermore, our experience with the ACME IBM/360-50 suggests that one could easily overestimate the ease of retaining compatibility, even within the IBM line, of programming systems that address distinctive objectives. (DOS systems differ from OS systems to a degree comparable to the barriers between machines from different manufacturers.) The best we can expect to do, in the face of conflicting objectives, is to strive for the most efficacious compatibility we can achieve in higher level languages, files systems, etc. We do have the benefit of Dr. John McCarthy's long experience with the PDP-line in the Artificial Intelligence Laboratory at the edge of the Stanford Campus.

These decisions have been reviewed by Dr. Gene Franklin in his role as University Associate Provost for Computing.

The SUMEX resource will divert its users from existing computer facilities only to a limited degree. Most of the uses intended for SUMEX require services simply not available otherwise. Some of the LISP programming of the DENDRAL project would otherwise be run on one of the larger IBM machines. This would be at great cost, and in any case could not permit an approach to closed loop management of the laboratory instruments. The entire ACME machine, scheduled with no shared users, lacks the computing power required for these applications.

We are recommending a twin PDP-10 for the SUMC service facility to optimize the overall advantages of the SUMEX option. IBM's performance in delivering time-sharing and realtime software for the 360 line has been disappointing -- indeed this was an important crimp in the projected transfer of routine time-sharing service from ACME to the SCC's 360-67 based service. On the other hand, DEC has done rather well with manufacturer-supplied software for these users of the PDP line.

An important advantage to the Medical Center of the dual PDP-10 system would be the availability of a back-up system. The lack of redundant hardware has precluded some applications; one user has opted to buy two mini-systems with identical configurations in order to obtain maximum reliability. The financial burden of this approach would be too great in most applications. The availability of redundant hardware facilities will be an important factor in the Hospital's consideration of how to solve its computing problems.

The impact on the Medical Center of having shared data files for research, service (research support), and administration can be significant. Faculty members in several disciplines are asking with increasing frequency for access to data bases outside of their own department. Of course, the file system design will provide file integrity, protection from catastrophic loss of data, and security. The availability of shared common files will (1) reduce the need for duplicate files, (2) improve the visibility and availability of information to faculty and staff, and (3) encourage placement of data on large, less expensive, rotating memories as opposed to smaller, more expensive hardware on satellite systems.

Common communications support for research service, and administrative computing systems is viewed as mandatory. The proliferation of terminal types, small machine types, etc. will breed a foul nest of communications hardware and software unless a sound, centralized, long-term plan is established. One can envision users with multiple terminals in each office, ward, or laboratory and multiple protocols needed by staff to use disparate systems unless the trends toward decentralization are encompassed under some umbrella of sound planning. The use of a shared communication system may permit redundancy and availability which would not be feasible in multiple independent systems. However, the hardware choices of the SUMC service operation are perhaps less important than the cooperative spirit that will be reinforced by the administrative arrangements for its coordination with SUMEX.

D. RESOURCE OPERATIONS

1. Administrative Structure

Line authority for SUMEX is vested in the P.I., Dr. Joshua Lederberg, who was also P.I. for the ACME system. He also functions as Chairman of the Department of Genetics, and, as a member of the Medical School's Executive Committee, is in good communications with the other department chairmen. In other roles (e.g., University Committees on Research and Computing Facilities, the Human Biology Program, etc.) he is also in frequent communication with general university activities. The Genetics Department includes the Instrumentation Research Laboratory directed by Dr. E. C. Levinthal.

Dr. Clayton Rich, as Dean, is the principal administrative officer for the Medical School, and Dr. Lederberg reports to him in several capacities.

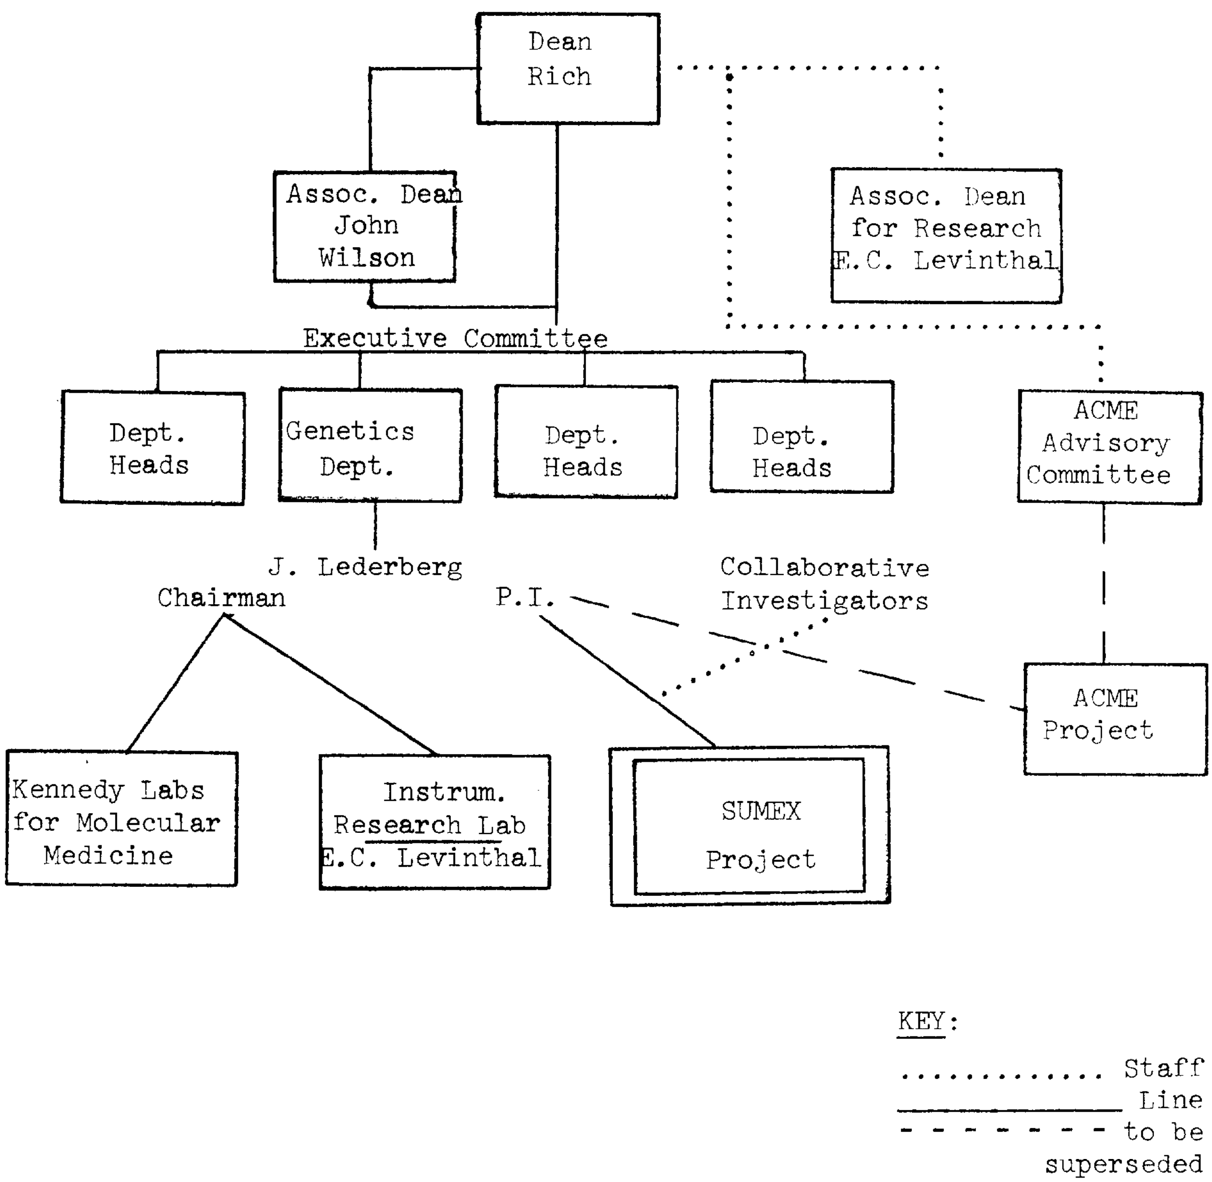

With the phasing out of ACME, its service responsibilities will be the responsibility of another organization reporting to the Dean independently of SUMEX. One proposed arrangement is:

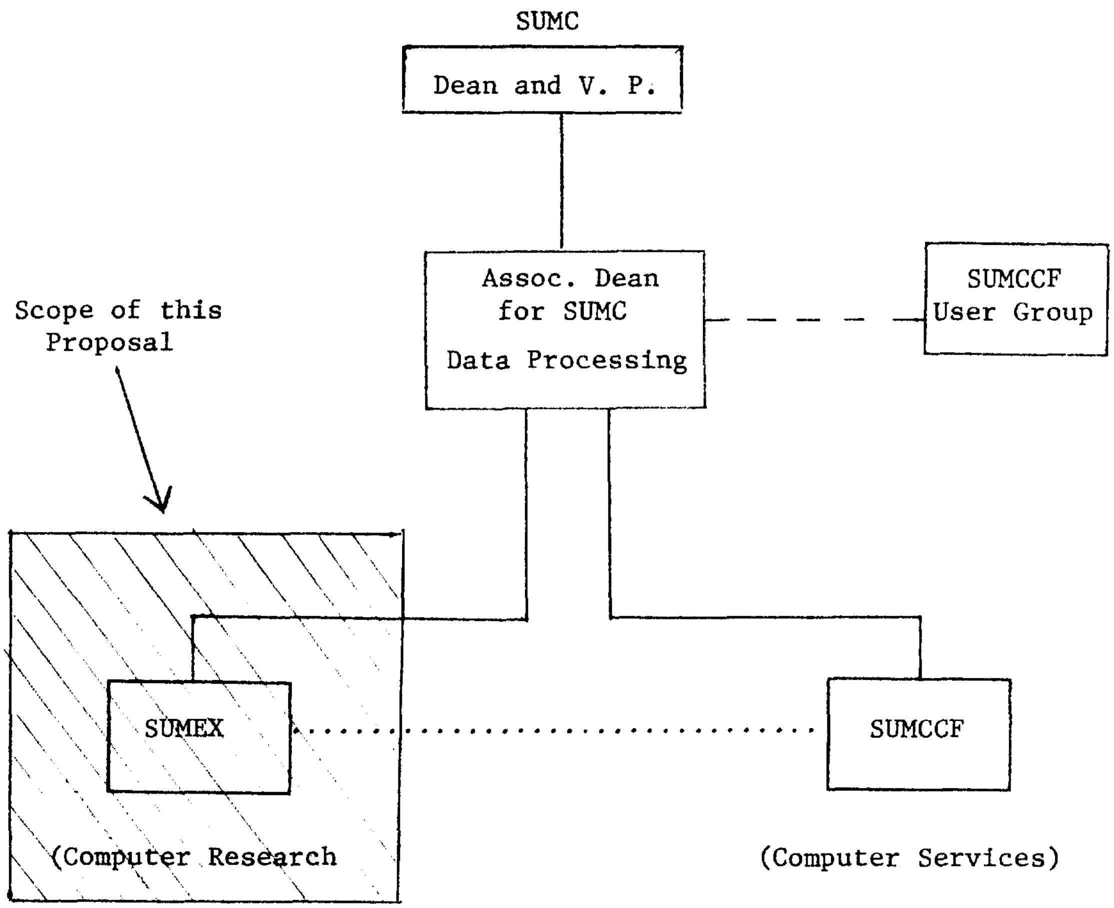

The Associate Dean will assure the optimum exchange of information and compatible policy development between computer research and computer services.

The hatched area signifies that this proposed grant will support (1) SUMEX operations and (2) liason with Stanford University Medical Center Computer Facility (SUMCCF). SUMEX will support development efforts on its PDP-10 machine of a kind that should be readily transferable to the SUMCCF as well. SUMEX will also assist the ACME community in the transition to SUMCCF services. However, operating expenses of the SUMCCF will be met from the SUMC budget including user fees, and not from the SUMEX grant. SUMEX and SUMCCF may however, provide mutual back up, on mutually advantageous terms, with respect to downtime emergencies and experiments involving linked processors, on the basis of credits for demand availability and use. The functions of the existing ACME facility will then be partitioned between SUMEX, for a limited number of computer-research oriented collaborations, and a combined computer service facility for the S. U. Medical Center, which we will label SUMCCF.

CURRENT

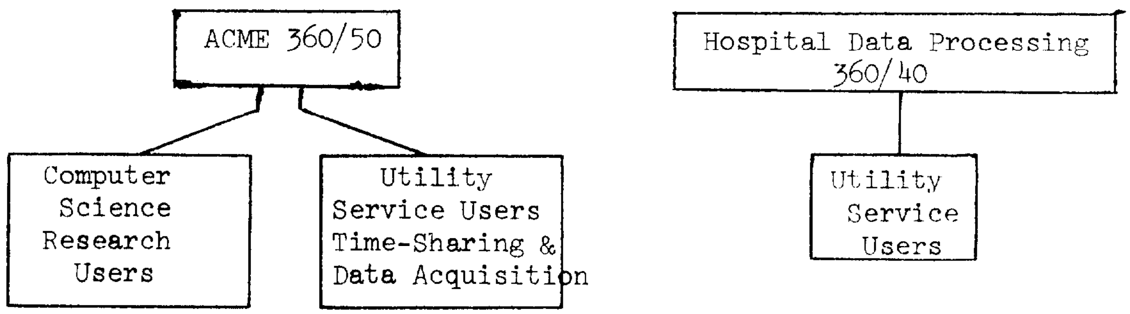

PROPOSED

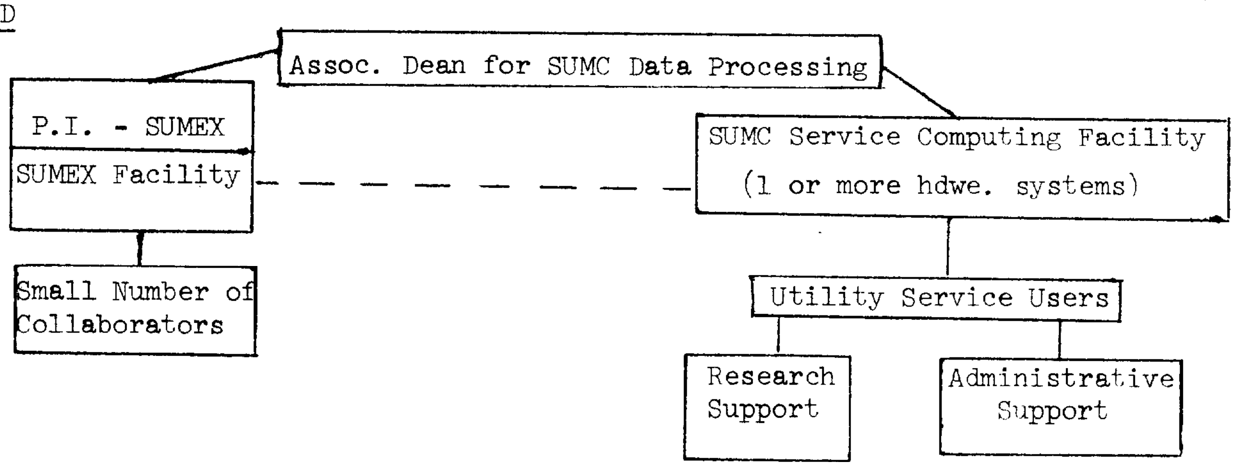

The SUMEX resource will operate under the direction of the Principal Investigator, who will be responsive to the research needs of the collaborators in terms of scheduling, use of the resources, and relative priorities for the programming staff. He will also establish liaison with SUMCCF to maximize the secondary gains of that relationship.

2. Operating Procedures and Policies

SUMEX is primarily responsive to a designated set of investigators interested in and competent to participate in major innovations in medical research applications of computers. They will have the opportunity to develop these applications in SUMEX prior to mounting them as new services on the Service Facility.

Authorized collaborators are limited and selected on the following grounds:

a. The Research Facility must remain capable of being dedicated to one experimenter's efforts if the total resource is needed for his work.

b. New hardware installation is likely to occur frequently on this facility. It should not have to be performed at odd hours in order to avoid normal service interruptions. The primary mission of this facility will be to service research users; routine service operations will be available only on the Service Facility.

c. Systems programmers will be testing new software concepts frequently. As a consequence, high system reliability and availability are not warranted.

d. Experience gained at ACME indicates that as time passes,more and more users come to expect (and demand) routine, stable, highly available service. Computer science related research cannot function efficiently in that environment.

e. Opening the Research Facility to routine use by many would inhibit the evolution of a fee-based Service Facility.

f. Management of the enterprise, including the selection of appropriate projects, will become increasingly difficult as the number of authorized users grows.

g. Their needs cannot be met on other local service facilities without undue disruption.

In the proposed grant period, no user service fees are contemplated for the Research Facility. As research concepts are developed and tested, it is expected that the Service Facility will be able to add services to meet user needs. Thus, the research objectives, once met, will be replaced with new research goals on the Research Facility.

Additional research collaborators will be recruited from the Stanford biomedical research community as indicated in paragraphs a. through g., above. We also contemplate cooperating with NIH grantees at other institutes via network facilities. During the interim period, prior to the settling down of reliable PDP-10 service on the SUMC facility, it may be desirable to coopt affiliated investigators who do not meet the full range of criteria but who are developing major projects in anticipation of the availability of the PDP-10 capabilities. Arrangements for servicing such users and for adding principal collaborators will be coordinated with the Biotechnology Research Resources Branch at regular intervals. The principal investigator will be responsible for applying these criteria for collaborating and affiliated projects, and for regular reporting to the Branch.

The Service Facility could provide routine services to SUMEX systems staff at times when the SUMEX was dedicated to a particular user's tasks. In addition, Service Facility usage would be needed to test newly transitioned packages from the Research Facility. The purchase of computer time on the SUMCCF would enhance the efficiency of SUMEX personnel. For these reasons, some funds are being requested to pay for services on the Service Facility.

E. COMPUTER CONFIGURATION RATIONALE

1. Introduction

This section addresses the problems of system configuration design and computer selection based on projected requirements and available machines. To summarize the discussions of these topics which follow, we have arrived at the following conclusions and proposed course of action:

(a) Separate machines for computer research and utility service are required to provide simultaneously a continuous and reliable computing utility service like the current ACME system and support for new system developments.

(b) The two machines, including required communications and future data base interfaces, ideally should be as similar as possible and geographically contiguous to allow redundancy for reliability, ease of software transfer, and efficiency of operation.

(c) The two machines will have some level of coordinated management but with financing of the service machine derived from fee for service funds and of the research machine from the presently proposed grant funds.

(d) Based on evolving software requirements for time-shared and realtime support as well as the capabilities and economics of currently available systems, it is proposed that both machines be Digital Equipment Corporation PDP-10 computers.

(e) A phased transition (Figure E-5) from the present single IBM 360/50 configuration to the dual PDP-10 configuration including necessary PL/ACME modifications is planned so as to minimize the trauma of conversion.

2. Computing Environment

The design of a medical experimental computing resource for research on satellite machine interactions and extended realtime problems interacts strongly with the overall design of computing support within the Stanford Medical Center. Based on past ACME experience, hospital administration experience, and projected Medical Center needs, an overall facility must be able to accommodate three main types of computing simultaneously:

(a) Medical Service Computing - A stable and reliable computing utility service must be available which supports on-going medical research and clinical needs in the sense that ACME currently performs these functions. The users of this type of utility are presently largely within Stanford but can be expected to extend outside of Stanford as network facilities come more and more into use. Such a utility must include in its repertoire appropriate state-of-the-art services for time-sharing and batch operation as well as satellite machine programming and on-line data communications facilities.

(b) Hospital Administrative Computing - A stable and reliable computing support of hospital administrative computing must be available for processing data related to patient accountability, financial records, clinical laboratory records, pharmacy records, etc. This type of computing is based to a considerable extent on software packages which have been developed outside of Stanford for specific computing systems. In the future the system must be able to support some level of integrated hospital information system. A long term requirement exists for compatible file structures accessible from various machines and software packages over local and larger scale networks.

(c) Medical Computing Research - Computing service must be available to support the development of computer system software and hardware capabilities as well as research projects which require sporadic dedication of large amounts of computing resources or which endanger system reliability. Such a service must be tolerant of higher system volatility than the utility services in order to allow evolutions in system design and utilization without impacting essential on-going computing functions.

There are reliability, capability, and priority conflicts in the requirements which these three groups place on a computing facility. The evolution of the present ACME system, while successful in making powerful computing tools easily and broadly available to medical researchers and clinicians, has also provided examples of such conflicts between various users. The ideal facility design must embed support for these various computing functions in an overall configuration which optimizes the desirable interactions of information and technology while minimizing the fundamental conflicts. As needs for computing resources within the Stanford Medical Center and its affiliates grow, the computing facility must be capable of economical expansion based on these needs in ways which minimize conversion and transition trauma.

3. Technical Requirements

Estimates of requirements for future computing service in terms of capacity, response time, communications, etc. are based on past experience with existing systems as well as projected new requirements. In the following only the requirements related to this grant application are considered. The major non-administrative computing service offered in the Medical Center has been the interactive, time-shared PL/ACME system. This type of system will continue to be the basic environment for the proposed research in satellite machine support and realtime systems. Thus the evolution of the ACME system is an essential element of this plan.

ACME Background - The PL/ACME system currently runs on an IBM 360/50 computer system shown functionally in Figure E-1. The time-sharing aspects of the system, developed under the previous ACME grant, apportion memory resources to multiple users from a large fixed reservoir (2.1 x 10^6 bytes). Realtime support for on-line experiments is provided by means of interfaces through either an IBM 1800 computer or an IBM 2701 data adapter. From a consideration of the loading history of this machine and related usage data a number of conclusions are drawn.

(a) The time-shared PL/ACME system has been of great benefit in fostering the medical use of computers at Stanford. It is expected that the needs for these services will increase in volume and sophistication.

(b) The 360/50 processor does not have the through-put capacity to provide adequate response service to existing heavy loads and is inadequate for closure of sophisticated realtime loops.

(c) Even with the large core memory available, core limitations impact performance and accessibility. A more sophisticated allocation of resources based on swapping and hardware relocation or on paging is required.

(d) The allocation of priorities to running tasks is too democratic with a resultant impact on applications with critical response timing requirements. Additional sophistication in the hardware and software priority hierarchy is necessary.

(e) Satellite machine programming and communication as well as real time needs will increase in terms of number of machines, complexity of application, aggregate data rates, and number of users. The system must integrate more flexible hardware and software satellite machine support into its repertoire.

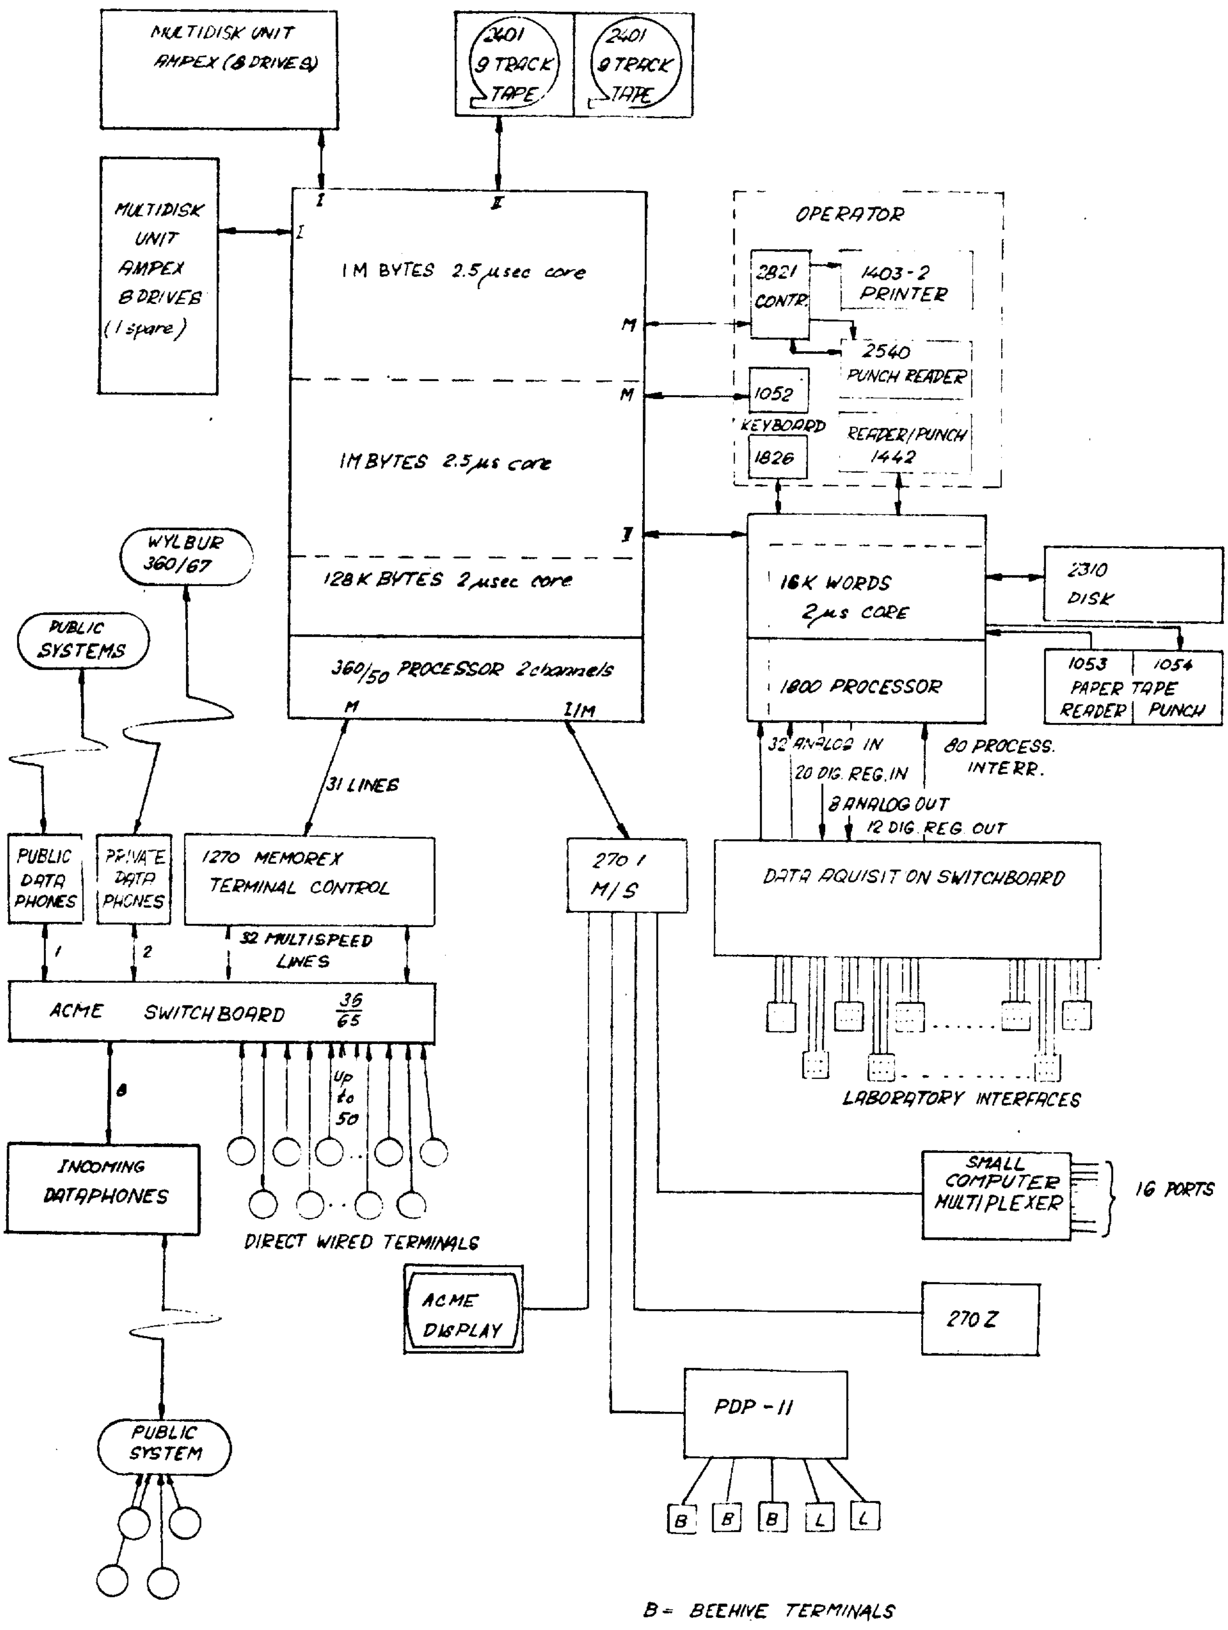

FIGURE E-1

ACME, May 1972

(f) More sophisticated usage of the system requires access to additional languages (such as list processing languages), more flexibility in organizing programs in terms of overlays, and more flexibility in overlapping task functions.

Future Needs - Based on these considerations, existing benchmarks for program through-put, and estimated new hardware and software requirements in support of research goals, the following gross summary of facility requirements is appropriate:

CPU: 2-4 times the through-put of the 360/50 with address relocation or paging and an interrupt hierarchy.

Memory: Approximately 10^6 bytes of memory.

Bulk Storage: 4-8 x 10^8 bytes.

User Load: Research (this grant): 20-40 terminals (4 to 6 projects).
Service: 50 - 100 terminals.

Satellite Machine Interfaces: Satellite machines will be used both for remote instrument interfaces and for local multiprocessing host support. Capabilities for direct memory sharing as well as normal satellite interfaces to the host as terminals are required.

System Software: The system software must allow effective scheduling and integration of conflicting time-sharing, batch and realtime computing loads.

4. Configuration Topologies

The requirements and priorities for computing services within the Medical Center place conflicting constraints on a computer facility. Some users require continuous, highly reliable service and others want access to develop system hardware and software capabilities which have an attendant risk of system crashes. Still others want sporadic dedication of sizable computing resources with response time constraints, and administrative elements of the computing load utilize software packages which require specific hardware and operating system characteristics (e.g. IBM developed business systems). One can consider a variety of facility configurations to attempt to meet these needs ranging from a large single processor to networks to completely distributed machines satisfying specific local needs.

Single Central Machine - The single machine topology is not a satisfactory solution because the conflicting priorities and reliability requirements cannot be resolved simultaneously in an adequate fashion. Serial scheduling of the conflicting usage is not feasible either because of time constraints. The duty cycle for reliable service support for medical research, clinical service, and hospital administration can be expected to approach 24 hours per day. The requirement for reliability implies more than one machine to guarantee minimum interruption in service.

Network Facilities - Machines of adequate capacity and with applicable software system capabilities exist at a number of nodes on established networks such as the ARPANET. Currently, however, because of the developmental nature of specific facilities or operational constraints, no known node or set of nodes can commit the needed capacity with suitable response time and reliability to meet estimated Medical Center needs. These factors rule out network facilities in the near term for providing computing services.

Distributed Local Computers - The completely distributed system is equally unsatisfactory at present because of gross inefficiency. Rarely can individual users keep a machine fully utilized for most research work proceeds sporadically. Whereas the cost of processors is decreasing dramatically, the cost of peripherals and memory as well as small machine programming is still high. Thus either gross inefficiency results or the individual researcher must make do with a machine of less capability than required. It is not feasible in the near future to completely rely on distributed machines without the necessary ability to synergize and focus large processing capability on specific problems as required.

Dual Central Machine Compromise - This leads to a compromise dual machine topology for the Medical Center computer facility which meets the simultaneous needs of conflicting users. The general characteristics of such a system are shown in Figure E-2. This system allows the maintenance of reliable computing service by shifting system element commitments at the expense of lower priority uses in the event of some failure. Conflicting requirements are resolved through the scheduling of separate processors. The Research and Service subsystems should be of similar configuration allowing interchangeability to achieve reliability for service support. In addition the symmetry of this configuration provides for software interchange in the evolution of the fruits of computing research to a more routine service environment. Economy is advanced by geographic contiguity implying shared operations as well as peripheral and communications equipment.

The service computing subsystem has as its basic design goals the distribution of reliable computing support for PL/ACME time-share users, realtime computing as it develops, small machine support, and administrative computing. This system will be funded on a fee for service basis using the currently projected ACME customer base as well as anticipated network users.

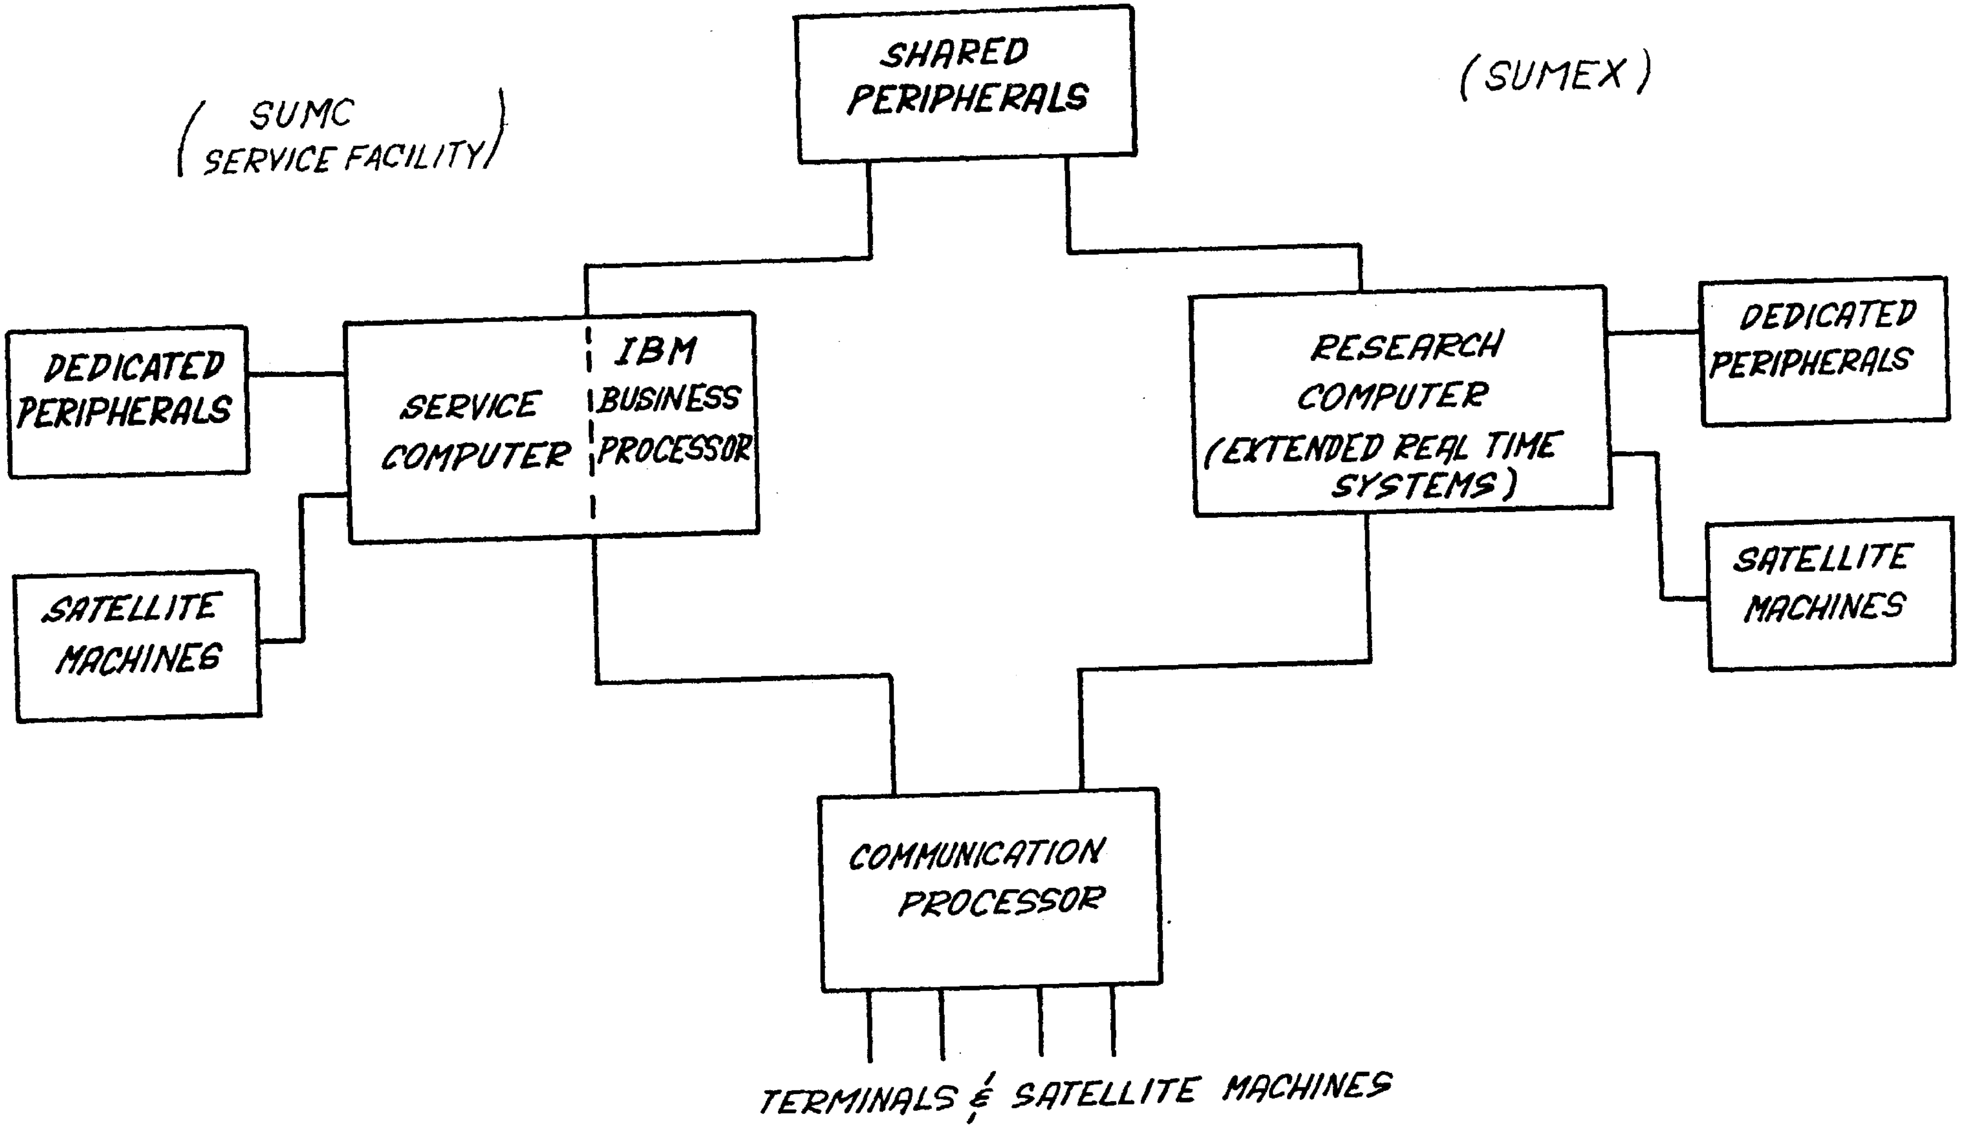

GENERAL CHARACTERISTICS OF THE PROPOSED SYSTEM

FIGURE E-2

The research computing subsystem has as its basic design goals the support of work under this grant proposal. These needs include an evolving PL/ACME based system to develop satellite computing support (remote and local) and extended realtime systems in an environment tolerant of developmental system volatility. This system also provides back-up to the service machine in case of failure and allows service system modification without impacting on-going service. The research system will be funded out of this grant if approved and will include provision for documenting and passing developed capabilities into the service domain as they become available.

5. Hardware Selection

a. Main Frame

The selection of hardware elements to carry out the facility plan shown in Figure E-2 requires consideration of several issues:

(a) The current configuration and its suitability or adaptability for future goals.

(b) Alternative hardware systems including capability and reliability.

(c) Alternative software systems including capability and reliability.

(d) User community contributions to software developments.

Current System - As indicated previously, the ACME system is an interactive time-sharing concept implemented on a large memory IBM 360/50 computer. The system as it exists supplies a powerful service to the Stanford community and could supply that service to a broader community given expanded through-put capacity. Such expansion in capacity is possible by incorporating a faster processor with more memory or a faster processor with a swapping relocation core sharing scheme. IBM will eventually make such features available on the more reliable 370 computer line. The current 512K byte memory limitation on the 370/145, however, is not contemplated to be expanded. This would force an expanded implementation in the much more expensive 370/155 or in an old machine, the 360/65-67.

Future Needs - The proposed research in this grant application, based on the PL/ACME system, places several additional requirements on the future system. First, the system must allow interfacing other languages (such as LISP, Assembly language, etc.) in addition to PL-1 and permit a more sophisticated hierarchical structuring of programs and their interfaces to on-going realtime events. This implies a system which at one level is as easy to use as the present PL/ACME system but which affords, when desired, the opportunity for deeper control of computer functions without having to convert to an entirely new system environment. It is recognized that these levels of user control may endanger system reliability through user-caused software crashes.

Secondly, the original system design did not foresee the proliferation of small machines which has occurred. Our research goals in the present plan include developing effective means for supporting and incorporating such systems in the distribution of computing services. This implies a requirement for more flexible ways of interfacing large and small machines in the form of memory to memory and CPU to external memory linkages. These go beyond the simple terminal or I/O device data forwarding roles currently in use.

Thirdly the design of realtime systems which allow "intelligent" loop closure commitments within time-sharing environments requires in addition to satellite machine interface flexibility, internal machine priority hierarchies, basically hardware implemented but with software control and extension.

Manufacturer Considerations - The size limitations of the 370/145, the high cost of the 370/155, the design age of the 360/65, and design goal considerations, suggest a reassessment of the selection of IBM hardware. This suggestion is enhanced by IBM's currently loose commitments to future time-sharing and realtime system support. The IBM systems provide a primitive priority and interrupt hierarchy and little flexibility in satellite processor interfaces.

From cost, hardware, and software points of view DEC equipment appears attractive. The PDP-10 hardware system (in particular the KI-10 processor) is of moderate capacity approximately comparable to the 360/65 and costs 30% less than the 370/155. The system is expandable to a multi-processor configuration to increase capacity at relatively low cost and has features comparable to those available on IBM hardware. These include soft fail machine check, hardware address relocation and paging (IBM will announce these shortly), and instruction look-ahead. In addition to these features, however, are a hierarchical interrupt structure and direct memory interfaces to satellite machines including other PDP-10's as well as PDP-11, and PDP-15 minicomputers. The software system is efficient and is designed around the integration of time-sharing, batch processing, and realtime computing taking advantage of appropriate hardware features. The PDP-10 community includes many research facilities working on aspects of machine architecture, system development, language support, artificial intelligence graphics, etc. These features of the PDP-10 and the computer science user community appear to be more in line with projected research and service requirements.

Other manufacturers offer hardware which incorporates many similar advantages but these systems lack the evolved software systems and applications packages as well as the large, actively working community of contributers which the DEC system has. Whereas IBM has a large community of users, the present impetus in the directions appropriate to this research are questionable. Thus, based on current information, it is felt that the PDP-10 offers a better technical and more

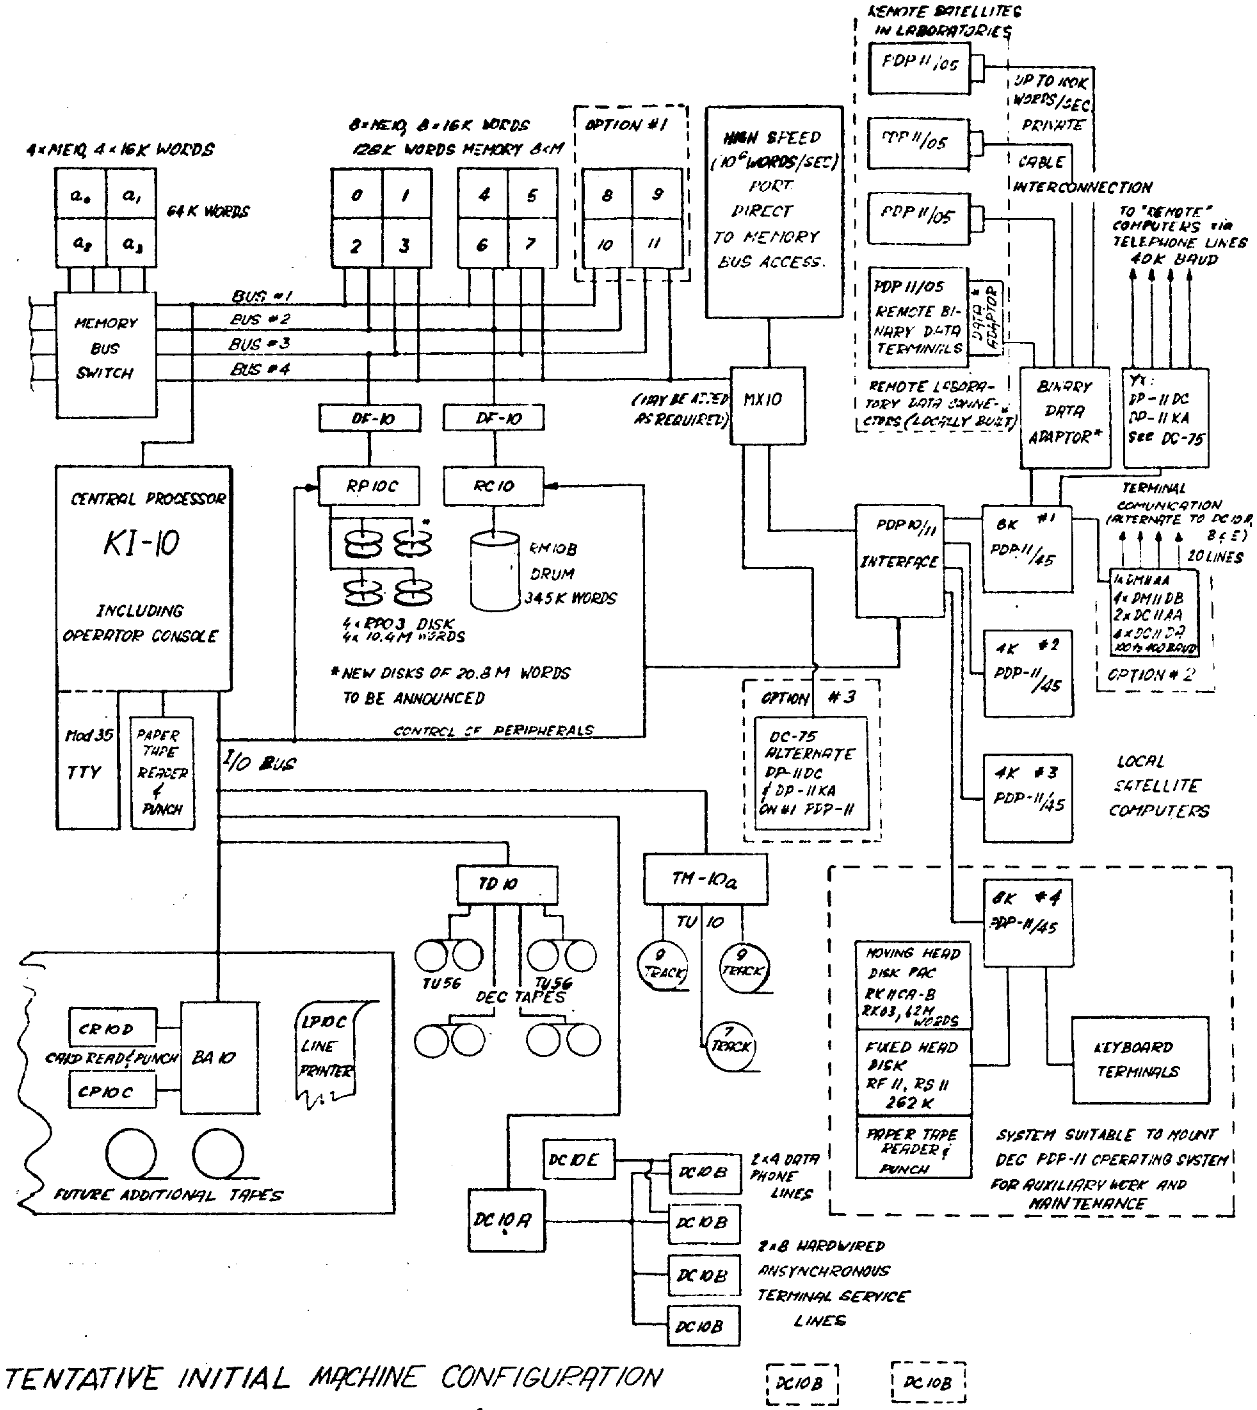

TENTATIVE INITIAL MACHINE CONFIGURATION

for

Stanford University Medical Center Experimental Computer Facility

"SUMEX"

*locally built

Figure E-3

economical long term posture for the experimental computing facility. This decision implies the judgment that the cost of converting the ACME software to run on a PDP-10 is offset by the longer term advantages expected to accrue from the more sophisticated system hardware and software support. The conversion effort is estimated to take 5.5 man years. A hardware change is painful but is less so the earlier it is made.

<u>Hardware Configuration</u> - The computer configuration as currently planned for the research subsystem is shown in Figure E-3. The service machine could appear as a symmetrical system.

b. Peripherals, Data Channels, and Satellite Computers

Attention is given here to the balance of the proposed hardware system. The RM10B drum is essential to DEC core swapping software. Certain other facilities are standard: The card reader and punch, line printer and operator's console need no explanation. The TD-10 controller and DEC tape units are necessary for maintenance of the PDP-10 system, hence a minimum configuration is included. 9 track tapes are selected for archive dumps, and some job entry or data interfaces with other computer facilities. A 7 track unit is included as experience has taught us that this tape format is still in use and a research facility must cope with this format from time-to-time.

<u>Data Channels</u> - This computer will quite literally exist for the processing of data from and to distant sites. Thus the remote data channels are of utmost interest. It is not intended to make a research project out of the communication aspects. On the contrary, it is proposed to have a set of solutions that can be implemented promptly and predictably in any laboratory, observing the due constraints of distance, cable, and rates. This service is to include diagnostic means for verification or troubleshooting.

Four classes of external communications are provided for:

1. Asynchronous character. 100 to approximately 400 baud. Suitable for TTY and other keyboard devices, many CRT terminals.

2. Synchronous character. Nominally up to 40,000 baud. Communication is suitable for telephone lines and will use nationally recognized standards. The modems used will be of standard commercial design, the proposed system uses exclusively DEC modems on the central computer end.

   The user may use commercial equipment of his choice on his end; however, there will be a unit available that incorporates a small remote satellite computer that is capable of running diagnostic programs to verify data communication to any remote location. Also this unit will be a "preferred" design for standardization.

3. Binary data channels. Up to 100K words/second. This would be restricted to on-campus connections that have multi-pair cable connections. Probably there would be a 3 pair, 11 pair and 19 pair version with somewhat different maximum speeds.

   The Binary Data Adaptors will be locally built. A characteristic is that both ends will plug into a PDP-11, or identical Unibus Models for this exist at Stanford in two versions, the IRL Chemistry-Medical Center connection, and the ACME small computer interface.

   Experience with many versions of these interfaces have taught the necessity of having standard service and connections to avoid the repetition of special engineering and resulting difficulty in maintaining service. The configuration of Figure E-4 provides standard service and is a configuration that may be checked by diagnostic software.

4. The need may arise for a superspeed data interface. None has been standardized for this purpose, but provision is being made to access Memory Bus #4 through the MX10 multiplexer for this purpose.

<u>Alternate Ports for Types 1, 2, or 3 Channels</u> - It is proposed currently to bring type 2 and type 3 channels into a local satellite computer. A trade-off exists here between hardware and software costs which will require further investigation. There is no commercial standard DEC service for type 2 to the memory bus or the I/O bus. For this and economy reasons the satellite PDP-11 is indicated. The hardware connection is to the PDP-11 bus, hence the configuration allows flexibility in moving any type 2 or type 3 channel to any of the local satellite computers.

<u>Satellite Computers</u> - To connect a PDP-11 to the PDP-10, the PDP-10/11 interface is indicated. This is a powerful but expensive device. It allows the PDP-11 to use segments of PDP-10 core. The segments allotted, and the interrupt service between computers, is enabled by the PDP-10. The PDP-10/11 interface also allows up to 8 PDP-11's; 4 are included in the present configuration. Incremental PDP-11's are economical and this allows exciting possibilities of small computer arrays for pipeline and parallel processing. This extra usefulness is thus a low priced expansion of the PDP-11/10 capability. The fourth PDP-11 has been configured as a DEC disc operating system. It can be used as a facility in a purely DEC fashion to assist other DEC users and the SUMEX staff in their software and interfacing efforts.

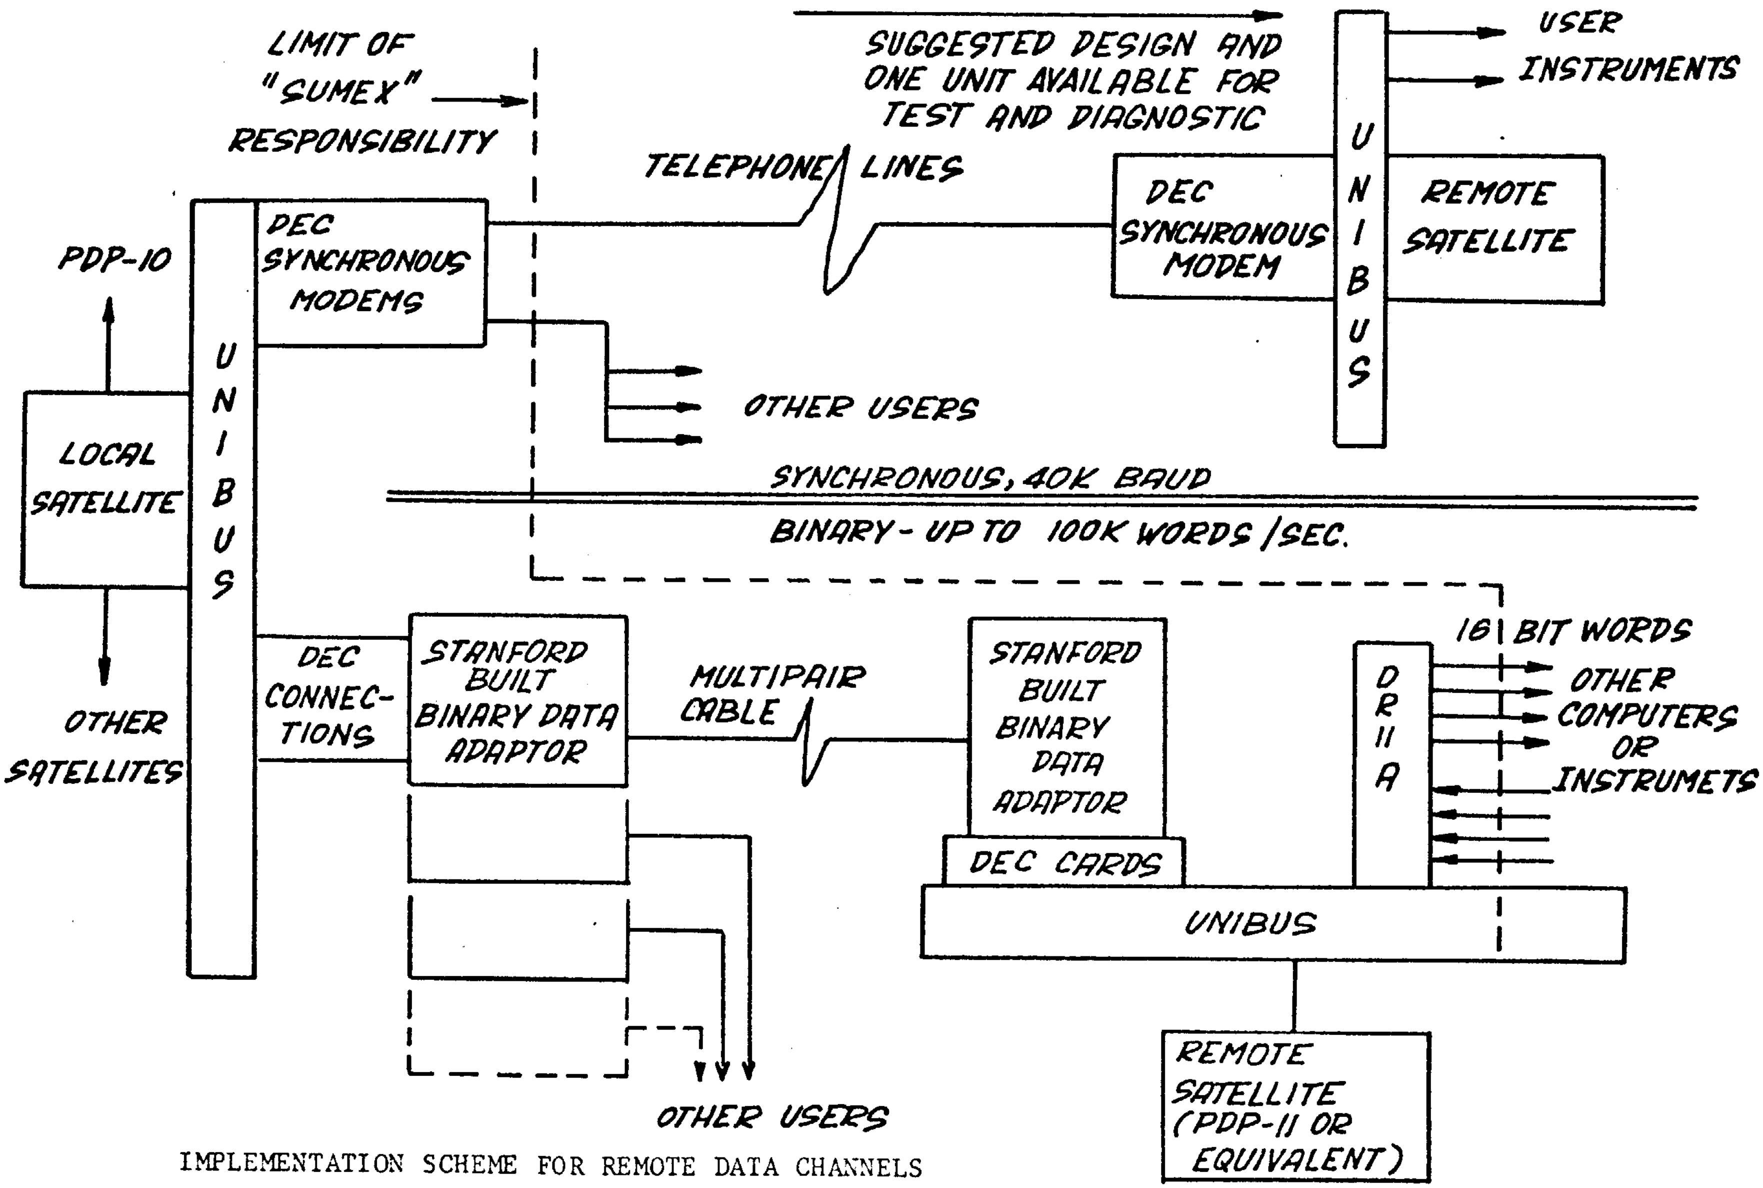

IMPLEMENTATION SCHEME FOR REMOTE DATA CHANNELS

FIGURE E-4